

Exhibit L



PATENTS
112056-0126
P01-1108

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re The Application of:)
Steven R. Kleiman et al.)
)
Serial No.: 10/105,034) Examiner: Nguyen, Than Vinh
)
Filed: March 21, 2002) Art Unit: 2187
)
For: A METHOD FOR WRITING)
CONTIGUOUS ARRAYS OF)
STRIPES IN A RAID STOR-)
AGE SYSTEM USING)
MAPPED BLOCK WRITES)
)
Cesari and McKenna, LLP
88 Black Falcon Avenue
Boston, MA 02210
February 22, 2005

"Express Mail" Mailing-Label Number: EV 433575215 US

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

PRELIMINARY AMENDMENT

This Preliminary Amendment is filed with a Request for Continued Examination filed on even date herewith in response to the Final Office Action mailed on November 19, 2004. All objections and rejections are respectfully traversed.

PATENTS
112056-0126
P01-1108

IN THE TITLE:

Please replace the previous title with the following new title:

- A METHOD FOR WRITING CONTIGUOUS ARRAYS OF STRIPES IN A
RAID STORAGE SYSTEM USING MAPPED BLOCK WRITES -

PATENTS
112056-0126
P01-1108

IN THE CLAIMS:

Please re-write the claims to read as follows:

- 1 1-16. (Cancelled)
- 2 17. (Currently Amended): A method for controlling storage of data in a plurality of storage
- 3 devices each comprising storage blocks configured in a plurality of parallel stripes, the
- 4 method comprising:
 - 5 buffering a plurality of write requests associated with data blocks for a single write
 - 6 transaction;
 - 7 defining a group of storage blocks, the group comprising a plurality of storage blocks
 - 8 in each of the plurality of storage devices; and
 - 9 associating mapping each data block with a respective one of the storage blocks
 - 10 across the plurality of stripes, for transmitting the association mapping to a storage device
 - 11 manager for processing of the single write transaction.
- 1 18. (Currently Amended): The method of claim 17, further comprising receiving by a file
- 2 system the plurality of write requests, and transmitting from the file system to the storage de-
- 3 vice manager the mapping association for processing of the single write transaction, wherein
- 4 the storage device manager comprises a RAID layer.
- 1 19. (Currently Amended): The method of claim 17, wherein mapping associating comprises
- 2 associating each data block of at least one of the write requests with storage blocks of only
- 3 one of the plurality of storage devices.

PATENTS
112056-0126
P01-1108

1 20. (Original): The method of claim 17, wherein defining the group of storage blocks com-
2 prises receiving by the file system from the storage device manager configuration informa-
3 tion of the plurality of storage devices.

1 21. (Original): The method of claim 20, wherein the configuration information comprises
2 disk topology information.

1 22. (Original): The method of claim 17, wherein the group of storage blocks comprises
2 more than one stripe in a RAID group.

1 23. (Original): The method of claim 22, wherein the group of storage blocks comprises
2 more than one contiguous stripe.

1 24. (Currently Amended): The method of claim 23, further comprising transmitting the
2 mapping association and the data to the plurality of storage devices to store each data block
3 at the respective one of the storage blocks of the more than one contiguous stripe. |

1 25. (Original): The method of claim 22, wherein the plurality of storage devices comprises
2 more than one RAID group, and defining the group of storage blocks comprises defining the
3 group of storage blocks from storage blocks in a first RAID group.

1 26. (Currently Amended): The method of claim 25, further comprising defining a second
2 group of storage blocks in a second RAID group for association with data associated-mapped
3 with a second plurality of write requests for processing of a second write transaction. |

PATENTS
112056-0126
P01-1108

1 27. (Currently Amended): The method of claim 17, further comprising queuing a plurality
2 of mappings associations by the storage device manager for processing of a plurality of write
3 transactions.

1 28. (Original): The method of claim 27, further comprising processing the plurality of write
2 transactions by the storage manager in an order derived from a plurality of priorities associ-
3 ated with the plurality of write transactions.

1 29. (Original): The method of claim 17, wherein the group of storage blocks comprises a
2 same quantity of storage blocks in each one of the plurality of storage devices.

1 30. (Original): The method of claim 29, wherein each of the plurality of storage blocks
2 comprises contiguous storage blocks.

1 31. (Original): The method of claim 17, wherein the group of storage blocks comprises allo-
2 cated and unallocated storage blocks.

1 32. (Original): The method of claim 31, wherein defining the group further comprises se-
2 lecting an unallocated storage block having a lowest block identifier of any unallocated stor-
3 age block as a storage block having a lowest storage block identifier of any storage block in
4 the group.

1 33. (Original): The method of claim 17, wherein the write requests comprise data to be
2 written.

PATENTS
112056-0126
P01-1108

- 1 34. (Currently Amended): A method for storing data blocks, the method comprising:
 - 2 providing a RAID layer in communication with a plurality of storage devices that
 - 3 each comprise a plurality of storage blocks configured in a plurality of parallel stripes;
 - 4 receiving by the RAID layer a write transaction request that includes an-a mapping
 - 5 ~~association~~ of each data block with a respective one of a group of storage blocks, the group
 - 6 comprising a plurality of storage blocks in each of the plurality of storage devices across the
 - 7 plurality of stripes; and
 - 8 storing the data blocks by the RAID layer in the group of storage blocks according to
 - 9 the mapping~~association~~.
- 1 35. (Original): The method of claim 34, wherein the RAID layer comprises a second plurality of storage devices that each comprise a plurality of storage blocks that mirror storage blocks of the plurality of storage devices, and further comprising storing the data in a group of storage blocks of the second plurality of storage devices.
- 1 36. (Original): The method of claim 35, further comprising performing a parity determination for one of the groups of storage blocks prior to storing, and sharing a result of the parity determination with the other of the groups of storage blocks.
- 1 37. (Original): The method of claim 36, further comprising selecting one of the groups of storage blocks for performing the parity determination.
- 1 38. (Original): The method of claim 34, further comprising monitoring by the storage device manager of configuration information of the plurality of storage devices.

PATENTS
112056-0126
P01-1108

1 39-41. (Cancelled)

1 42. (Currently Amended): A device for controlling storage of data in a plurality of storage
2 devices each comprising storage blocks configured in a plurality of parallel stripes, the de-
3 vice comprising:

4 a buffer collecting write requests, each request associated with data blocks;
5 a processor defining a group of storage blocks, the group comprising a plurality of
6 storage blocks in each of the plurality of storage devices, the processor mapping associating
7 each data block with a respective one of the storage blocks across the plurality of stripes for a
8 single write transaction; and
9 a transmitter transmitting the mapping association to a storage device manager for
10 processing of the single write transaction.

1 43. (Currently Amended): An apparatus for storing data blocks, the apparatus comprising:

2 a plurality of storage devices that each comprise a plurality of storage blocks config-
3 ured in a plurality of parallel stripes; and
4 a storage device manager in communication with the plurality of storage devices, and
5 configured to receive a write transaction from a file system, the write transaction comprising
6 a mapping an association of each data block to a respective storage block of a group of stor-
7 age blocks, the group comprising a plurality of storage blocks in each of the plurality of stor-
8 age devices across the plurality of stripes.

PATENTS
112056-0126
P01-1108

1 44. (Original): The apparatus of claim 43, wherein the storage device manager causes a par-
2 ity determination for the group of storage blocks prior to storing the data blocks in the group
3 of storage blocks.

1 45. (Currently Amended): A device for controlling storage of data in a plurality of storage
2 devices each comprising storage blocks configured in a plurality of parallel stripes, the de-
3 vice comprising:

4 means for buffering a plurality of write requests associated with data blocks for a sin-
5 gle write transaction;

6 means for defining a group of storage blocks, the group comprising a plurality of
7 storage blocks in each of the plurality of storage devices; and

8 means for mapping associating each data block with a respective one of the storage
9 blocks across the plurality of stripes, for transmitting the mapping association to a storage
10 device manager for processing of the single write transaction.

1 46. (Currently Amended): A computer readable media, comprising: the computer readable
2 media containing instructions for execution in a processor for the practice of the method of,

3 buffering a plurality of write requests associated with data blocks for a single write
4 transaction;

5 defining a group of storage blocks, the group comprising a plurality of storage blocks
6 configured in a plurality of parallel stripes in each of a plurality of storage devices; and

7 mapping associating each data block with a respective one of the storage blocks
8 across the plurality of stripes, for transmitting the mapping association to a storage device
9 manager for processing of the single write transaction.

PATENTS
112056-0126
P01-1108

1 47. (Currently Amended): Electromagnetic signals propagating on a computer network,
2 comprising: the electromagnetic signals carrying instructions for execution in a processor for
3 the practice of the method of,

4 buffering a plurality of write requests associated with data blocks for a single write
5 transaction;

6 defining a group of storage blocks, the group comprising a plurality of storage blocks
7 configured in a plurality of parallel stripes in each of a plurality of storage devices; and

8 mapping associating each data block with a respective one of the storage blocks
9 across the plurality of stripes, for transmitting the mapping association to a storage device
10 manager for processing of the single write transaction.

1 48. (Currently Amended): An apparatus for storing data blocks, the apparatus comprising:

2 means for providing a RAID layer in communication with a plurality of storage de-
3 vices that each comprise a plurality of storage blocks configured in a plurality of parallel
4 stripes;

5 means for receiving by the RAID layer a write transaction request that includes a
6 mapping an association of each data block with a respective one of a group of storage blocks,
7 the group comprising a plurality of storage blocks in each of the plurality of storage devices
8 across the plurality of stripes; and

9 means for storing the data blocks by the RAID layer in the group of storage blocks
10 according to the mapping association.

1 49. (Currently Amended): A computer readable media for use with storing data blocks,
2 comprising: the computer readable media containing instructions for execution in a processor
3 for the practice of the method of,

PATENTS
112056-0126
P01-1108

4 providing a RAID layer in communication with a plurality of storage devices that
5 each comprise a plurality of storage blocks configured in a plurality of parallel stripes;

6 receiving by the RAID layer a write transaction request that includes a mapping an
7 association of each data block with a respective one of a group of storage blocks, the group
8 comprising a plurality of storage blocks in each of the plurality of storage devices across the
9 plurality of stripes; and

10 storing the data blocks by the RAID layer in the group of storage blocks according to
11 the mapping association.

1 50. (Currently Amended): Electromagnetic signals propagating on a computer network for
2 use with storing data blocks, comprising: the electromagnetic signals carrying instructions for
3 execution in a processor for the practice of the method of,

4 providing a RAID layer in communication with a plurality of storage devices that
5 each comprise a plurality of storage blocks configured in a plurality of parallel stripes;

6 receiving by the RAID layer a write transaction request that includes a mapping an
7 association of each data block with a respective one of a group of storage blocks, the group
8 comprising a plurality of storage blocks in each of the plurality of storage devices across the
9 plurality of stripes; and

10 storing the data blocks by the RAID layer in the group of storage blocks according to
11 the mapping association.

1 51. (Currently Amended): A method for controlling storage of data, comprising:

2 receiving one or more write requests associated with data blocks;

3 receiving topological information associated with storage blocks configured in a plu-
4 rality of parallel stripes of a storage system;

PATENTS
112056-0126
P01-1108

5 mapping associating the data blocks with one or more storage blocks across the plu-
6 rality of stripes; and

7 writing the data blocks, in response to the mapping association, to the one or more
8 storage devices in a single write request.

1 52. (Currently Amended): The method of claim 51, further comprising: transmitting the
2 mapping association to a storage device manager.

1 53. (Currently Amended): The method of claim 51, further comprising: mapping associating
2 each data block with a single storage block.

1 54. (Currently Amended): The method of claim 51, further comprising: storing the data
2 blocks in the mapping association.

1 55. (Previously Presented): The method of claim 51, further comprising: storing the data
2 blocks in a memory of the storage system.

1 56. (Currently Amended): The method of claim 51, further comprising: creating an array as
2 the mapping association.

1 57. (Previously Presented): The method of claim 51, further comprising: buffering a plural-
2 ity of write requests into the single write request.

PATENTS
112056-0126
P01-1108

1 58. (Previously Presented): The method of claim 57, further comprising: buffering, in a
2 buffer, the plurality of write requests into the single write request until a predetermined crite-
3 ria is met.

1 59. (Previously Presented): The method of claim 58, further comprising: meeting the criteria
2 when a buffer is full.

1 60. (Previously Presented): The method of claim 58, further comprising: meeting the criteria
2 when the single write request is a predetermined logical length.

1 61. (Previously Presented): The method of claim 51, further comprising: receiving topologi-
2 cal information of a plurality of storage devices of the storage system.

1 62. (Previously Presented): The method of claim 61, further comprising: using a RAID
2 system as the plurality of storage devices.

1 63. (Previously Presented): The method of claim 51, further comprising: using a plurality of
2 disks for the storage system.

1 64. (Currently Amended): The method of claim 51, further comprising: transmitting the
2 mapping association to the storage device manager.

1 65. (Currently Amended): The method of claim 51, further comprising: organizing the map-
2 ping association as a combination of columns and rows.

PATENTS
112056-0126
P01-1108

1 66. (Currently Amended): The method of claim 65, further comprising: mapping associating
2 each column with a storage device.

1 67. (Currently Amended): The method of claim 66, further comprising: mapping associating
2 each row with a stripe of the storage device.

1 68. (Currently Amended): The method of claim 67, further comprising: arranging the map-
2 ping association as a plurality of stripes on a plurality of storage devices.

1 69. (Currently Amended): A storage system, comprising:

2 a file system, the file system to receive one or more write requests associated with
3 data blocks;

4 a storage device manager, the storage device manager to generate topological infor-
5 mation of storage blocks configured in a plurality of parallel stripes of one or more storage
6 devices, and to send the topological information to the file system; and

7 a mapping an-association generated in the file system, the mapping association to as-
8 sociate map the data blocks with one or more storage blocks across the plurality of stripes of
9 the one or more storage devices, the mapping association to be sent to the storage device
10 manager, the storage device manager to write the data blocks, in response to the mappingas-
11 sociation, to the one or more storage blocks as a single write request.

1 70. (Currently Amended): The storage system of claim 69, further comprising: a mapping
2 an-association that associates maps each data block with a single storage block.

PATENTS
112056-0126
P01-1108

1 71. (Previously Presented): The storage system of claim 69, further comprising: a memory
2 to buffer the data blocks for the write request to the one or more storage devices.

1 72. (Currently Amended): The storage system of claim 69, further comprising: a memory to
2 store the mapping association containing the data blocks.

1 73. (Previously Presented): The storage system of claim 69, further comprising: one or more
2 storage devices having storage blocks.

1 74. (Currently Amended): The storage system of claim 69, further comprising: an array as
2 the mapping association.

1 75. (Previously Presented): The storage system of claim 69, further comprising: a buffer in
2 the file system to receive the one or more write requests.

1 76. (Currently Amended): The storage system of claim 69, further comprising: an arrange-
2 ment of columns and rows in the mapping association.

1 77. (Previously Presented): The storage system of claim 76, further comprising: each col-
2 umn representing a storage device of the one or more storage devices.

1 78. (Previously Presented): The storage system of claim 76, further comprising: each row
2 representing a stripe of the one or more storage devices.

PATENTS
112056-0126
P01-1108

1 79. (Previously Presented): The storage system of claim 69, further comprising: a plurality
2 of disks as the one or more storage devices.

1 80. (Previously Presented): The storage system of claim 69, further comprising: a RAID
2 system as the plurality of storage devices.

1 81. (Currently Amended): A storage system, comprising:

2 means for receiving one or more write requests associated with data blocks;

3 means for receiving topological information associated with storage blocks config-
4 ured in a plurality of parallel stripes of a storage system;

5 means for mapping associating the data blocks with one or more storage blocks
6 across the plurality of stripes; and

7 means for writing the data blocks, in response to the mapping association, to the one
8 or more storage devices in a single write request.

1 82. (Currently Amended): A computer readable media, comprising: the computer readable
2 media containing instructions for execution in a processor for the practice of the method of,

3 receiving one or more write requests associated with data blocks;

4 receiving topological information associated with storage blocks configured in a plu-
5 rality of parallel stripes of a storage system;

6 mapping associating the data blocks with one or more storage blocks across the plu-
7 rality of stripes; and

8 writing the data blocks, in response to the mapping association, to the one or more
9 storage devices in a single write request.

PATENTS
112056-0126
P01-1108

- 1 83. (Currently Amended): Electromagnetic signals propagating on a computer network,
- 2 comprising: the electromagnetic signals carrying instructions for execution in a processor for
- 3 the practice of the method of,
- 4 receiving one or more write requests associated with data blocks;
- 5 receiving topological information associated with storage blocks configured in a plu-
- 6 rality of parallel stripes of a storage system;
- 7 mapping associating the data blocks with one or more storage blocks across the plu-
- 8 rality of stripes; and
- 9 writing the data blocks, in response to the mappingassociation, to the one or more
- 10 storage devices in a single write request.

PATENTS
112056-0126
P01-1108

REMARKS

This Preliminary Amendment is filed with a Request for Continued Examination filed on even date herewith in response to the Final Office Action mailed on November 19, 2004. All objections and rejections are respectfully traversed.

Claims 17-38, and 42-83 are in the case.

Claims 17-19, 24, 26-27, 34, 42-43, 45-54, 56, 64-70, 72, 74, 76, and 81-83 have been added to better claim the invention.

At paragraph 3 of the Office Action, the title of the invention was objected to as being non-descriptive. A new title has been provided, and the title is believed to be in allowable condition.

At paragraph 4 of the Office Action, Examiner points out that the reference “DAVID HITZ et al” on the IDS filed 5/22/03 was not considered because it lacks a date. Applicant has filed herewith a new IDS with the date of publication for the Hitz reference, namely, March 1995. Applicant directs Examiner to the section entitled “Preface to the Reader,” which states, “This paper reflects the nomenclature and product characteristics in March 1995.” Accordingly, Applicant respectfully requests that the “DAVID HITZ et al” reference be considered prior art for this current application.

At paragraph 5 of the Office Action, claims 47, 50, and 83 were rejected under 35 U.S.C. §101 as being directed to non-statutory subject matter on the grounds that electromagnetic signals set out in these claims are not tangibly embodied.

PATENTS
112056-0126
P01-1108

The present invention, as set out in representative claim 47, comprises in part:

47. Electromagnetic signals propagating on a computer network, comprising: the electromagnetic signals carrying *instructions for execution in a processor for the practice of the method of,*

buffering a plurality of write requests associated with data blocks for a single write transaction;

defining a group of storage blocks, the group comprising a plurality of storage blocks configured in a plurality of parallel stripes in each of a plurality of storage devices; and

mapping each data block with a respective one of the storage blocks across the plurality of stripes, for transmitting the mapping to a storage device manager for processing of the single write transaction.

Applicant respectfully urges that the novel method steps are tangibly embodied in the electromagnetic signals propagating on the computer network. Further, Applicant respectfully urges that the embodiment of electromagnetic signals for transfer of *instructions for execution in a processor for the practice of the method of* between computers fully satisfies all requirements of 35 U.S.C. 101, and all requirements set out in the MPEP.

That is, Applicant respectfully urges that embodiment of the instructions in electromagnetic signals meets all of the requirements of 35 U.S.C. § 101, especially as clarified by MPEP 2106 IV, B, 1(c) at page 2106 of MPEP 8th Edition Incorporating Revision No. 2. (hereinafter MPEP 2106 IV, B, 1(c)). Further, MPEP 2106 IV, B, 1(c) states, at page 2106:

“However, a signal claim directed to a practical application of electromagnetic energy is statutory regardless of its transitory nature, see *O'Reilly* 56 U.S. at 114-19; *In re Breslow*, 616 F. 2d 516, 519-21, 205 USPQ 221, 225-26 (CCPA 1980).”

In the case *In re Breslow* claims were permitted by the Court (CCPA) to chemical species which are transient in nature, and cannot be separated out of the mixture in which

PATENTS
112056-0126
P01-1108

they are created. The MPEP cites this patentability of transitory phenomena in chemical reactions in support of the statement by the MPEP, "However, a signal claim directed to a practical application of electromagnetic energy is statutory regardless of its transitory nature."

The important point for patentability is the practical application of electromagnetic energy. And a practical application of electromagnetic energy is transmission of a computer program over a computer network, where the computer program is for the practice of a novel method. This practical application of electromagnetic energy is patentable subject matter, as explained by MPEP 2106 IV, B, 1(c).

A copy of *In re Breslow* from 205 USPQ 221 is attached to this Amendment, for the convenience of the Examiner.

Applicant respectfully urges that imbedding instructions for execution on a processor in an electromagnetic signal propagating on a computer network meets the practical application requirements of 35 U.S.C. § 101 and of MPEP 2106 IV, B, 1(c), and that claim 47 therefore claims statutory subject matter. Also, Applicant respectfully urges that claims 50 and 83 claim statutory subject matter under 35 U.S.C. § 101 and MPEP 2106 IV, B, 1(c).

At paragraph 8 of the Office Action, claims 17-38, and 42-44 were rejected under 35 U.S.C. §102 (b) as being anticipated by DeKoning, U.S. Patent No. 6,148,368, issued on November 14, 2000.

The present invention, as set forth in representative claim 17, comprises in part:

17. A method for controlling storage of data in a plurality of storage devices each comprising storage blocks configured in a plurality of parallel stripes, the method comprising:

buffering a plurality of write requests associated with data blocks for a single write transaction;

PATENTS
112056-0126
P01-1108

defining a group of storage blocks, the group comprising a plurality of storage blocks in each of the plurality of storage devices; and

mapping each data block with a respective one of the storage blocks across the plurality of stripes, for transmitting the mapping to a storage device manager for processing of the single write transaction.

DeKoning discloses a method for accelerating disk array write operations using segmented cache memory and data logging. In DeKoning, write commands are first buffered to the cache segments, and later accumulated into a cache extension disk region at a later time, such as when the disks are idle. The accumulated write requests are the write commands combined into “larger write requests (e.g. RAID stripe writes)” (Col. 8, Lines 4-7, emphasis added) to spare the storage system from the write penalty of multiple smaller writes. Notably, the accumulated write request is a single stripe write (Col. 2, Lines 21-25).

Applicant respectfully urges that DeKoning does not show Applicant’s claimed novel, “*mapping each data block with a respective one of the storage blocks across the plurality of stripes, for transmitting the mapping to a storage device manager for processing of the single write transaction.*”

Applicant’s claimed invention is directed toward the mapping of data blocks with the storage blocks to which they will be written. Applicant claims buffering a plurality of write requests and combining them into a single write transaction. In addition to this, however, Applicant goes one step further by mapping each data block of the single write request with a storage block across a *plurality of stripes* of the storage system before transmitting the buffered write request to a storage device manager. In this way, the storage device manager is not required to map each data block to a storage block of the storage system across the *plurality of stripes*, as it would have conventionally been required to do. Particularly, Applicant’s claimed invention is not limited to sending single stripe

PATENTS
112056-0126
P01-1108

write transactions as DeKoning is, but sends write requests as mapped (associated) blocks spanning multiple stripes of the storage devices.

Applicant respectfully points out the following excerpts from the Specification of the present invention to further explain the novel claimed invention:

In a preferred embodiment, the invention features a file system that supports *coordinated writes to multiple stripes*. The file system receives write requests and disk topology information. The file system creates the block layout information responsive to the disk topology information and the write request, and thus prepares a single I/O write transaction for processing by a RAID layer. (Page 16, Line 29 – Page 17, Line 3, *emphasis added*)

And further:

The block layout information 15 is forwarded to the disk array manager 13 for processing of the write transaction. The disk array manager 13 receives the write transaction request, which includes the block layout information 15 that identifies one or more stripes for the single write transaction. *Prior art systems, in contrast, typically send single stripe write transactions to a RAID layer. Features of the invention thus enable forwarding of a complete write transaction to a RAID layer.* (Page 19, Lines 19-27, *emphasis added*)

DeKoning does not address mapping data blocks with storage blocks across a plurality of stripes for transferring a single write transaction, but instead merely discusses the use of buffering smaller write requests into a larger write request, such as a RAID (single) stripe write.

Notably, DeKoning states in its Background that:

It is advantageous to accumulate small data I/O requests in cache memory so that the aggregation of their data may be written to the disk array as a single "stripe write" (an operation involving all disks of the array)

PATENTS
112056-0126
P01-1108

in parallel operation). Striping large data I/O requests across many disks allows parallel access allows quick reading or writing of large amounts of data. If a write-back-cache is present, then writes can be deferred, thus minimizing the write Penalty. With a large write-back cache, data can accumulate in cache and be consolidated into a single write so that only one read-modify-write sequence and disk seek and rotation operation needs to be done for multiple consolidated writes. (Col. 2, Lines 21-33, emphasis added)

DeKoning continues this discussion by suggesting that there is a need to solve various problems involved with management of the cache memory.

Examiner cites to the paragraphs contained between Col. 7, Line 45 and Col. 8, Line 13 of the DeKoning patent as a basis for rejection under 35 U.S.C. §102. Particularly, the DeKoning patent states therein that: “The logged information includes the write data as well as metadata identifying the destination of the write data in the LUN to which it was addressed” (Col. 7, Lines 48-51).

However, as Applicant pointed out above, the DeKoning patent is directed to managing cache memories to combine multiple write requests into larger RAID (single) stripe writes. DeKoning clearly is solving a different problem than Applicant’s claimed invention, as pointed out by the following statement in DeKoning:

Furthermore, storing write requests in cache-extension disk region 232, an extension of the log structure in segments 235, 236, and 238,, as opposed to storing it in main disk region 234, allows further deferring of writing to main disk region 234 and the RAID write penalties associated therewith. Thus complexities and latencies of mapping data to main disk region 234, which may be configured as a RAID level 5, may be deferred until quieter periods. (Col. 10, Lines 30-38, emphasis added)

In summary, DeKoning combines multiple write requests into a single ordinary stripe write request. In sharp contrast, Applicant claims a further step, mapping the multiple write requests with the actual storage blocks across a plurality of stripes on the disks to which they are to be written.

PATENTS
112056-0126
P01-1108

Applicant respectfully urges that the DeKoning patent is legally precluded from anticipating the claimed invention under 35 U.S.C. §102 because of the absence from the DeKoning patent of Applicant's "*mapping each data block with a respective one of the storage blocks across the plurality of stripes, for transmitting the mapping to a storage device manager for processing of the single write transaction.*"

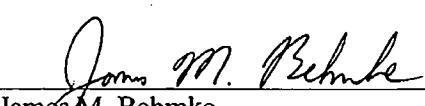
All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims, and therefore in condition for allowance.

Favorable action is respectfully solicited.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,


James M. Behmke
Reg. No. 51,448
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500